Congestion Avoidance Commands

This chapter describes commands used to avoid congestion.

bandwidth (QoS)

To specify the minimum bandwidth allocated to a class belonging to a policy map, use the **bandwidth** command in policy map class configuration mode. To remove the bandwidth specified for a class, use the **no** form of this command.

bandwidth {*rate* [*units*]|**percent** *percentage-value*}
no bandwidth {*rate* [*units*]|**percent** *percentage-value*}

Syntax Description

rate	Minimum bandwidth, in the units specified, to be assigned to the class. Range is from 1 to 4294967295.
units	Specifies the units for the bandwidth.
percent *percentage-value*	Specifies the amount of minimum guaranteed bandwidth, based on an absolute percentage of available bandwidth. Range is from 1 to 100.

Command Default

The default units is kbps.

Command Modes

Policy map class configuration

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

The **bandwidth** command is used to specify the minimum guaranteed bandwidth allocated for traffic matching a particular class. Bandwidth may be defined as a specific value or may be set as a percentage of the interface bandwidth.

If a percentage value is set, the accuracy that can be expected is 1 percent.

The **bandwidth** command is supported only in the egress direction.

A policy map can have a single bandwidth statement per class. Both percentage and actual value bandwidth configurations can be used within a policy map.

The **bandwidth** command does not specify how the bandwidth is to be shared. Instead it specifies how much bandwidth is guaranteed per class, by setting the number of tokens that are assigned to the token bucket of a particular class. For configured behavior to work correctly, you must ensure that the sum of the bandwidths plus any priority traffic is not greater than the bandwidth of the interface itself. If the interface is oversubscribed, unpredictable behavior results.

Task ID

Task ID	Operations
qos	read, write

Examples

This example shows how to guarantee 50 percent of the interface bandwidth to a class called class1 and 10 percent of the interface bandwidth to a class called class2:

```
RP/0/RP0:hostname(config)# policy-map policy1
RP/0/RP0:hostname(config-pmap)# class class1
RP/0/RP0:hostname(config-pmap-c)# bandwidth percent 50
RP/0/RP0:hostname(config-pmap-c)# exit
RP/0/RP0:hostname(config-pmap)# class class2
RP/0/RP0:hostname(config-pmap-c)# bandwidth percent 10
```

bandwidth remaining

To specify how to allocate leftover bandwidth to various classes, use the **bandwidth remaining** command in policy map class configuration mode. To return to the system defaults, use the **no** form of this command.

bandwidth remaining [{**percent** *percentage-value*|**ratio** *ratio-value*}]
no bandwidth remaining [{**percent** *percentage-value*|**ratio** *ratio-value*}]

Syntax Description

percent *percentage-value*	Specifies the amount of guaranteed bandwidth, based on an absolute percentage of the available bandwidth. Range is from 1 to 100.
ratio *ratio-value*	Specifies the amount of guaranteed bandwidth, based on a bandwidth ratio value. Range is 1 to 2000.

Command Default

No bandwidth is specified.

Command Modes

Policy map class configuration

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

Bandwidth, bandwidth remaining, shaping, queue-limit and WRED commands can be configured together in the same class.

Note The **bandwidth remaining** command is supported only in the egress direction.

The available bandwidth is equally distributed among those queueing classes that do not have the remaining bandwidth explicitly configured.

The **bandwidth remaining** command is used to proportionally allocate bandwidth to the particular classes, but there is no reserved bandwidth capacity.

On egress, if the **bandwidth remaining** command is not present, then the bandwidth is shared equally among the configured queueing classes present in the policy-map.

Task ID

Task ID	Operations
qos	read, write

Examples

This example shows how the remaining bandwidth is shared by classes class1 and class2 in a 20:80 ratio.

```
RP/0/RP0:hostname(config)# policy-map policy1
RP/0/RP0:hostname(config-pmap)# class class1
RP/0/RP0:hostname(config-pmap-c)# bandwidth remaining percent 20
RP/0/RP0:hostname(config-pmap-c)# exit
RP/0/RP0:hostname(config-pmap)# class class2
RP/0/RP0:hostname(config-pmap-c)# bandwidth remaining percent 80
```

queue-limit

To specify or modify the maximum number of packets the queue can hold for a class policy configured in a policy map, use the **queue-limit** command in policy map class configuration mode. To remove the queue packet limit from a class, use the **no** form of this command.

queue-limit *value* [*unit*]
no queue-limit

Syntax Description

value	Maximum threshold for tail drop in bytes. Range is from 1 to 4294967295.
unit	(Optional) Units for the queue limit value. **Note** When the specified *units* is packets, packets are assumed to be 256 bytes in size.

Command Default

The default value is 10 milliseconds for all queues including the high-priority queues.

Command Modes

Policy map class configuration

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

When configuring the **queue-limit** command, you must configure one of the following commands: **priority**, **shape average**, **bandwidth** or **bandwidth remaining**, except for the default class. The default value is 10 milliseconds for all queues including the high-priority queues.

The **queue-limit** command is supported only in the egress direction.

Packets satisfying the match criteria for a class accumulate in the queue reserved for the class until they are serviced by the scheduling mechanism. The **queue-limit** command defines the maximum threshold for a class. When that threshold is reached, enqueued packets to the class queue result in tail drop (packet drop). Tail drop is a congestion avoidance technique that drops packets when an output queue is full, until congestion is eliminated.

Use the **show qos interface** command to display the queue limit and other QoS values.

These restrictions apply to queue limits:

- Queue limit should be at least the maximum MTU size, which is 9 * 1024 bytes = 9kb.
- Queue limit should be 3 GB, which is the maximum packet buffer size in ingress and egress queueing ASICs.
- Only time-based units are allowed on bundle targets.

Guaranteed Service Rate

The guaranteed service rate is defined as the service rate of the queue when all queues are backlogged and derived as:

minimum_bandwidth + (bandwidth_remaining_percent * unallocated_bandwidth)

This example shows the guaranteed service rate calculation:

```
policy-map sample_policy
  class c1
    bandwidth percent 30
    bandwidth remaining percent 40
  class c2
    bandwidth percent 20
  class class-default
```

guaranteed service rate of c1 = 30 percent LR + (40 percent * 50 percent * LR)

guaranteed service rate of c2 = 20 percent LR + (30 percent * 50 percent * LR)

guaranteed service rate of class-default = 30 percent * 50 percent * LR

- Where LR is line rate of the target on which service policy "sample_policy" is attached.
- 50 percent is unallocated bandwidth.

Task ID

Task ID	Operations
qos	read, write

Examples

This example shows how to set the queue limit for a class to 1000000 packets for policy map policy1:

```
RP/0/RP0:hostname(config)# policy-map policy1
RP/0/RP0:hostname(config-pmap)# class class1
RP/0/RP0:hostname(config-pmap-c)# queue-limit 1000000
```

random-detect

To enable random early detection (RED), use the **random-detect** command in policy map class configuration mode. To remove RED, use the **no** form of this command.

random-detect {**default**|**discard-class** *value*|*min-threshold* [*units*] *max-threshold* [*units*]}
no random-detect

Syntax Description

default	Enables RED with default minimum and maximum thresholds.
discard-class *value*	Discard-class based WRED (up to 2 values, which is 0 and 1).
min-threshold	Minimum threshold in number of packets. The value range of this argument is from 0 to 1073741823 in packets.
max-threshold	Maximum threshold in number of packets. The value range of this argument is from the value of the *min-threshold* argument to 1073741823. When the queue length exceeds the maximum threshold, RED drops all packets with the specified discard class value.
units	(Optional) Units for the threshold values.

Command Default

Default unit for *max-threshold* and *min-threshold* is **packets**.

Command Modes

Policy map class configuration

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

The RED congestion avoidance technique takes advantage of the congestion control mechanism of TCP. By randomly dropping packets before periods of high congestion, RED tells the packet source to decrease its transmission rate. Assuming the packet source is using TCP, it decreases its transmission rate until all the packets reach their destination, indicating that the congestion is cleared. You can use RED as a way to cause TCP to slow transmission of packets. TCP not only pauses, but it also restarts quickly and adapts its transmission rate to the rate that the network can support.

RED distributes losses in time and maintains normally low queue depth while absorbing traffic bursts. When enabled on an interface, RED begins dropping packets when congestion occurs at a rate you select during configuration.

When time units are used, the guaranteed service rate is used to compute thresholds. The default minimum threshold is 6 ms and the maximum threshold is 10 ms.

When the value of the *units* argument is packets, packets are assumed to be 256 bytes in size.

Weighted Random Early Detection

The following restriction apply to Weighted Random Early Detection (WRED):

- For thresholds in time units, the guaranteed service rate is used to calculate the thresholds in bytes.

For bundles, queue limit and WRED thresholds are supported in time units only.

Task ID

Task ID	Operations
qos	read, write

Examples

This example shows how to enable RED using a minimum threshold value of 1000000 and a maximum threshold value of 2000000:

```
RP/0/RP0:hostname(config)# policy-map policy1
RP/0/RP0:hostname(config-pmap)# class class1
RP/0/RP0:hostname(config-pmap-c)# random-detect 1000000 2000000
```

service-policy (interface)

To attach a policy map to an input interface or output interface to be used as the service policy for that interface, use the **service-policy** command in the appropriate configuration mode. To remove a service policy from an input or output interface, use the **no** form of the command.

service-policy {**input**|**output**} *policy-map*
no service-policy {**input**|**output**} *policy-map*

Syntax Description

input	Attaches the specified policy map to the input interface.
output	Attaches the specified policy map to the output interface.
policy-map	Name of a service policy map (created using the **policy-map** command) to be attached.

Command Default

No service policy is specified.

Command Modes

Interface configuration.

Command History

Release	Modification
Release 6.1.42	This command was introduced.

Usage Guidelines

You can attach a single policy map to one or more interfaces to specify the service policy for those interfaces. The class policies composing the policy map are then applied to packets that satisfy the class map match criteria for the class. To apply a new policy to an interface, you must remove the previous policy. A new policy cannot replace an existing policy.

Task ID

Task ID	Operations
qos	read, write

Examples

This example shows policy map policy2 applied to HundredGigabitEthernet 0/0/0/1.

```
RP/0/RP0:hostname(config)# class-map class2
RP/0/RP0:hostname(config)# match precedence ipv4 2
RP/0/RP0:hostname(config-cmap)# exit

RP/0/RP0:hostname(config)# policy-map policy2
RP/0/RP0:hostname(config-pmap)# class-map class2
RP/0/RP0:hostname(config-pmap-c)# set precedence 3
RP/0/RP0:hostname(config-pmap)# exit

RP/0/RP0:hostname(config)# HundredGigabitEthernet 0/0/0/1
RP/0/RP0:hostname(config-if)# service-policy input policy2
```

This example shows policy map policy 1 applied to Bundle-Ether interface.

```
RP/0/RP0:hostname(config)# interface Bundle-Ether1
RP/0/RP0:hostname(config-if)# service-policy input policy1
RP/0/RP0:hostname(config-if)# exit
```